arXiv:1701.04941v1 [cs.SY] 18 Jan 2017

Chapter 1
Ensemble Control of Cycling Energy Loads: Markov Decision Approach

Michael Chertkov and Vladimir Y. Chernyak

Abstract Markov Decision Process (MDP) framework is adopted to represent ensemble control of devices, whose energy consumption pattern is of a cycling type, e.g. thermostatically controlled loads. Specifically we utilize and develop the class of MDP models, coined previously linearly-solvable MDP, that describe optimal dynamics of the probability distribution of an ensemble of many cycling devices. Two principally different settings are discussed. First, we consider optimal strategy of the ensemble aggregator aimed at balancing between minimizing the common objective that consists of the cost of the ensemble consumption, under varying price of electricity, and the welfare deviation from normal, evaluated as the KL-divergence between the optimal and nominal transition probabilities between different states of the ensemble's participants. Second, we shift to the Demand Response setting, where the aggregator aims to minimize the welfare deviation under the condition that the aggregated consumption of the ensemble matches the perfectly varying in time target consumption, requested by the system operator. We discuss a modification of both settings, aimed at encouraging or constraining the transitions between different states. The dynamic programming feature of the resulting modified MDPs is always preserved, however lin(ear)-solvability is lost fully or partially, depending on the type of modification. We also conducted some (limited in scope) numerical experimentation using the formulations of the first setting. We conclude with discussing future generalizations and applications.

Michael Chertkov
Theoretical Division, T-4 & CNLS, Los Alamos National Laboratory Los Alamos, NM 87545, USA and Energy System Center, Skoltech, Moscow, 143026, Russia, e-mail: chertkov@lanl.gov

Vladimir Y. Chernyak
Department of Chemistry, Wayne State University, 5101 Cass Ave, Detroit, MI 48202, USA, e-mail: chernyak@chem.wayne.edu

1.1 General Motivation/Introduction

Power systems, as well as other energy systems, undergo a transition from traditional device-oriented and deterministic approach, to a variety of novel approaches accounting for

- stochasticity and uncertainty in how the devices, especially new devices like wind-farms, are operated;
- network aspects, e.g., in terms of optimization, control and design/planning;
- utilizing massive amount of newly available measurements/data for the aforementioned optimization, control, and design settings.

Demand Response (DR) became one important component of the smart grid development, see e.g. [6, 33] and references therein. Novel DR architecture brakes the traditional paradigm, where only generators are flexible, suggesting that participation of loads in control can be beneficial for the grid at large, and also sufficiently comfortable for consumers/loads.

DR assumes that the loads follow operational commands of the System Operator (SO). Using DR to control the frequency in the range from tens of seconds to minutes is one such potentially very attractive niche for the use of the DR [32, 13, 3, 38]. Frequency control is a tool to balance production and consumption. Traditionally the frequency control is achieved by adjusting the generators, while DR (when developed) helps to achieve the balance by also adjusting the loads. The task of frequency control is set by the SO for loads as a temporal consumption request, which can be formulated ahead of time (say, for the next 10 min) or in real time, expecting frequent updates (say, every ten seconds).

This type of frequency control, namely DR services, are already provided by big consumers like aluminium smelters [35] and desalination plants [2]. However, the potential effect would be an order of magnitude larger if small loads are available to provide DR services. However, direct involvement of small-scale consumers is expensive due to associated communication and control costs. A potentially viable way out of this problem is to control many small-scale consumers indirectly – via an intermediate entity, also called the aggregator of an ensemble, including many (say thousands or tens of thousands) small-scale consumers [11, 22, 12, 27, 28, 25, 26, 7, 6, 4]. Therefore, this novel DR architecture assumes, action-wise, separation into the following three levels:

(a) SO sends control request/signal to an aggregator. The SO request is stated as a temporal profile (possibly binned into intervals, of say 15s) expected from the ensemble for an upcoming interval duration of, say, 10 min.

(b) Aggregator processes the SO signal by solving an optimization/control problem and communicating the same A-signal, which is the output/solution of the optimization/control problem, to all members of the ensemble (end-point consumers). The A-signal is broadcasted to the end-point devices/consumers, thus making the communication part cheap.

(c) End-point devices receive and implement the A-signal. Implementation is assumed to be straightforward, or at least requiring only light computations at a device for turning the A-signal into controllable actions.

There is a number of challenges, associated with this novel aggregator-based DR architecture. The challenges are of both conceptual/formulation and of the solution/implementation type. Let us mention a couple of these challenges, most relevant to level (b) of the control architecture which is in the focus of this manuscript.

In posing the A-optimization one is looking for a simple/minimal but also a sufficiently accurate way of modeling individual devices. In particular, one needs to decide on how to model the devices in terms of their achievable states and spatial-temporal resolution. To address the conceptual challenge we will state the aggregator-level problems utilizing the language of the Markov Decision Process (MDP), therefore controlling the transition probabilities between states. More accurately, we will assume that each device can find itself in all (or subset of all) of the finite number of states. Then we describe the probabilistic state of the ensemble in terms of the vector of non-negative numbers summed to unity, each representing the fraction of the consumers, observed in a given state at any given moment of time. The optimization/control degree of freedom for the A-optimization is represented by the set of stochastic transition probability matrices between the states, defined at each time slot over the discretized and finite time horizon.

In addition to the frequency control formulation we will also discuss in this manuscript another setting which is of interest by itself, but which is also less challenging in terms of finding efficient computational solution. In this algorithmically less complex case we will be looking for an optimal balance between the cost of the ensemble operations, when the price of electricity changes in time, and the deviation of the probabilistic state of the ensemble from what would be a natural/normal behavior of the ensemble in the case without the price bias.

The discrete-time, discrete-space modeling has its continuous-time, mixed-space counterpart known under the name of Thermostatically-Controlled-Loads (TCL), or even more generally cycling loads, characteristic of periodic (or quasi-periodic) evolution in the phase space. See [6] and our own recent paper [9] for detailed discussions of the TCL modeling and many relevant references. In this manuscript we choose to work with the discrete-time, discrete-space MP models for the sake of universality and flexibility of the latter. Indeed, an MP model can be viewed as a macro-model following from its TCL, micro-model counterpart in the result of a model reduction/coarse-graining in space-time, see e.g. [31]. However, the MPs represent much broader class of models, compared to those obtained via coarse-graining of the TCL counterparts, or even when compared to a combination (average) of different TCL models. This most general class of MDP models is especially useful in the context of Machine Learning (ML) where an MP is reconstructed from actual measurements/samplng of the underlying ensemble.

From MP, which represents stochastic dynamics, we transition to MDP that represents stochastic optimization with the MP/ensemble state constrained by the requirement of being within the class of state dynamics, described by the chosen class of MPs. Many choices of MDP would be reasonable from the prospectives

of practical engineering. In this manuscript we adopt and develop an approach, pioneered in [19], and further developed in [23, 24, 36, 16, 17, 30]. This relatively unexplored formulation of the stochastic optimal control is known under the names of the "path integral control" in the case of continuous-time, continuous-space formulation [23], and also as the "Linearly Solvable MDP" (LS-MDP) in the case of discrete-time, discrete-space setting [36, 16, 17]. Linear solvability of the simplest version of the scheme is advantageous, as it allows to carry out analytical expressions for the optimal solution one step further than done otherwise through the Dynamic Programming (DP) approach. (We remind that DP involves recursive solution the Bellman-Hamilton-Jacobi (BHJ) equations.) One extra advantage of the approach, which holds even in the case when reduction of the BHJ-equation to a linear equation is not possible, is related to the fact that the penalty term associated with deviation of the optimized/actual transition probability from its ideal shape (not perturbed by the SO signal) has a very natural form of the generalized Kublack-Leibler (KL) comparison of the two probability measures.

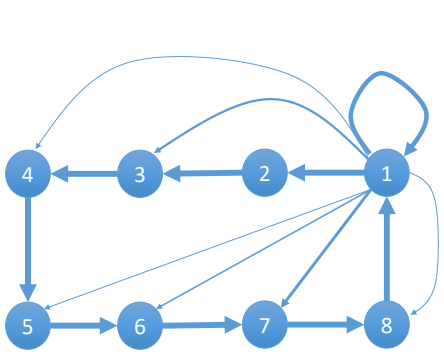

switch on | switch off

0.2	0.5	0.1	0.03	0.02	0.03	0.1	0.02
0.02	0.2	0.5	0.1	0.03	0.02	0.03	0.1
0.1	0.02	0.2	0.5	0.1	0.03	0.02	0.03
0.03	0.1	0.02	0.2	0.5	0.1	0.03	0.02
0.02	0.03	0.1	0.02	0.2	0.5	0.1	0.03
0.03	0.02	0.03	0.1	0.02	0.2	0.5	0.1
0.1	0.03	0.02	0.03	0.1	0.02	0.2	0.5
0.5	0.1	0.03	0.02	0.03	0.1	0.02	0.2

Fig. 1.1 An exemplary MDP resulting from space-binning/discretization and time-discretization of a two-level mixed-state TCL model, e.g. of the type discussed in [9]. This model is used in Section 1.4 for illustrative experiments. The directed graph of the "natural" transitions and respective stochastic $\bar{p}$ matrix are shown in the left and right Subfigures respectively.

When discussing MPs and MDPs we will utilize the following (rather standard) notations and terminology:

- Dynamic state of the ensemble is described in terms of a vector, $\rho(t) = (\rho_\alpha(t)|\forall\alpha)$, where $\rho_\alpha(t) \geq 0$ is the probability to find a device (from the ensemble) in the state α at time t. The vector is normalized (no probability loss) at all times considered, $\sum_\alpha \rho_\alpha(t) = 1, \quad \forall t, \quad \forall\alpha$.

- $p(t) = (p_{\alpha\beta}(t)|\forall t, \forall \alpha, \beta)$ is the transition probability matrix, describing an MP and set to be an optimization variable within the MDP framework. The matrix element, $p_{\alpha\beta}(t)$, represents the transition probability, i.e. the probability for a device which was at the state α at the moment of time t transition to the state β at $t+1$. The matrix is stochastic, i.e.

$$\sum_{\alpha} p_{\alpha\beta}(t) = 1, \quad \forall t = 0, \cdots, T-1, \quad \forall \beta, \tag{1.1}$$

- The evolution of the MP/ensemble state is described by the Master Equation (ME)

$$\rho_{\alpha}(t+1) = \sum_{\beta} p_{\alpha\beta}(t)\rho_{\beta}(t), \quad \forall t \; \forall \alpha. \tag{1.2}$$

  The ME should be supplemented by the initial condition for the state of the ensemble:

$$\rho_{\alpha}(0) = \rho_{in;\alpha}, \quad \forall \alpha \tag{1.3}$$

  Given the initial condition (1.3) one needs to run ME (1.2) forward in time to find the states of MP ensemble at all future times. The direction of time is principal, as it reflects physical causality of the setting.
- MDP formulations discussed in the manuscript are stated as optimizations over the matrix p. All of the formulations depend on the target $\bar{p}$, correspondent to p which would be optimal if the ensemble is "left alone" for sufficiently long time. In the case of the aforementioned frequency control the formulation discussed in Section 1.5, $\bar{p}$ corresponds to the steady state of the ensemble when the frequency tracking guidance is ignored. In the formulations of Sections 1.2, 1.3, $\bar{p}$ corresponds to the steady state of the ensemble, operating for sufficiently long time in the case of a flat, i.e. time and state independent, cost. An exemplary MP and respective "natural" $\bar{p}$, introduced in the result of coarse-graining of a TCL, is shown in Fig. (1.1).

The material in the remainder of the manuscript is organized as follows. MDP aimed at finding the profit optimal transition probability vector for an ensemble of devices is defined and analyzed in Section 1.2. MDP discussed in this Section measure the deviation of the "optimal" p, from its "normal" counterpart, $\bar{p}$, in terms of the standard KL divergence. This formulation is thus an LS-MDP, where we have placed the material related to a detailed technical discussion of the MDP's linear solvability and related features and properties in Appendix 1.9. We further discuss in Section 1.3 a modified MDP that differentiates between the transitions, i.e. discounts some transitions and encourages others, in comparison with differentiation-neutral formulation of the preceding Section. In Section 1.4 we experiment with the optimization objectives introduced and discussed in Sections 1.2,1.3. In Section 1.5 we describe and discuss the solution of an MDP, seeking to minimize the welfare deviation, stated in terms of the KL distance (or weighted KL-distance) between

the "optimal" and "normal" transition probabilities, while also matching the SO-objective exactly. Preliminary discussion of the MDP approach integration into the Optimum Power Flow, e.g., seeking to control the voltage and minimize the power losses, is given in Section 1.6. Section 1.7 is reserved for the conclusions and discussions of the path forward.

1.2 Profit (balanced with welfare) Optimal

Our main MDP formulation of interest, which we coin profit vs welfare, is stated as the following optimization

$$\min_{p,\rho} \mathbb{E}_\rho \left[\sum_{t=0}^{T-1} \sum_{\alpha} \left(\underbrace{U_\alpha(t+1)}_{\text{Cost of Electricity}} + \underbrace{\sum_{\beta} \log \frac{p_{\alpha\beta}(t)}{\bar{p}_{\alpha\beta}}}_{\text{welfare penalty}} \right) \right]_{\text{Eqs. (1.1,1.2,1.3)}} \tag{1.4}$$

$$= \min_{p,\rho} \sum_{t=0}^{T-1} \sum_{\alpha} \left(\left(U_\alpha(t+1) + \sum_{\beta} \log \frac{p_{\alpha\beta}(t)}{\bar{p}_{\alpha\beta}} \right) \rho_\alpha(t) \right)_{\text{Eqs. (1.1,1.2,1.3)}} , \tag{1.5}$$

where the matrix p is the optimization/control variable, which is stochastic, according to Eq. (1.1). Here in Eq. (1.4) $\bar{p}$ is an exogenously known stochastic matrix describing the transition probabilities, correspondent to normal dynamics/mixing within the ensemble, i.e. explaining the dynamics which the ensemble would show in the case of normal (welfare comfortable), $U = 0$, objective. The stochasticity of $\bar{p}$ means that $\bar{p}$ satisfies Eqs. (1.1), where p is replaced by $\bar{p}$.

We assume that when following $\bar{p}$ the ensemble reaches a statistical steady state, $\rho_\alpha^{(\mathrm{st})}$, i.e. mixes, sufficiently fast, $\bar{p}\rho^{(\mathrm{st})} = \rho^{(\mathrm{st})}$.

The essence of the optimization (1.4) is a compromise an aggregator aims to achieve in balancing between saving most for the ensemble while also keeping the level of discomfort (welfare penalty) to its minimum. The two conflicting objectives are represented by the two terms in Eq. (1.4).

Optimization (1.4) is rather general, while for the example of a specific discrete-time-discrete-space TCL, one sets U_α to non-zero only for the states α that represent the "switch-on" states of the TCL.

Eq. (1.4) is of the so-called linearly-solvable class of MDPs, introduced in [36], discussed in [16],[17], [29], and also briefly described, as a special case, in Appendix 1.9. Solution of Eqs. (1.4) is fully described by Eqs. (1.24,1.27,1.28).

According to the general description part of the Appendix 1.9, the problem is solved in two DP steps

- Backward in time step. Compute p recursively advancing backwards in time according to Eqs. (1.24,1.27) where $\gamma(t)$ is substituted by 1, and starting from the final condition (1.28).
- Forward in time step. Reconstruct ρ running Eqs. (1.2) forward in time with the initial condition (1.3).

1.3 Differentiating states through a penalty/encouragment

The penalized Kublack-Leibler welfare penalty term in Eq. (1.4) is restrictive in terms of how the transitions between different states, and also observed at different moments of time, compare to each other. To encourage/discourage, or generally differentiate the transitions one may weight the terms in the KL-sum differently, thus introducing the $\gamma_{\alpha\beta}(t)$ factors:

$$\min_{p,\rho}\mathbb{E}_\rho\left[\sum_{t=0}^{T-1}\sum_{\alpha}\left(\underbrace{U_\alpha(t+1)}_{\text{Cost of Electricity}}+\underbrace{\sum_{\beta}\gamma_{\alpha\beta}(t)\log\frac{p_{\alpha\beta}(t)}{\bar{p}_{\alpha\beta}}}_{\text{welfare penalty}}\right)\right]_{\text{Eqs. (1.1,1.2,1.3)}} \tag{1.6}$$

This generalization of Eq. (1.4) aims to ease the implementation of the optimal decision, e.g. emphasizing or downplaying the controllability of a transitions between the states and at different moments of time.

Solution of the optimization (1.6) is described in Appendix 1.9. Even though linear solvability of the state-uniform formulation, discussed in the preceding Section 1.2, does not extend to the nonuniform formulation of Eq. (1.6), the basic DP approach still holds and the problem is solved via the following backward-forward algorithm

- Backward in time step. Starting with the final conditions (1.17), one solves Eqs. (1.18) for φ recursively backwards in time. Solving Eqs. (1.18) requires, at each time step, to execute an inner loop iterations (till convergence), according to Eqs. (1.22), to find the Lagrangian multiplier λ. At each step backwards in time p is reconstructed according to Eqs. (1.19).
- Forward in time step. Reconstruct ρ running Eqs. (1.2) forward in time with the initial condition (1.3).

1.4 Computational Experiments

In this Section we describe some (preliminary) computational experiments, conducted for MDP settings discussed in the two preceding Sections. We choose for

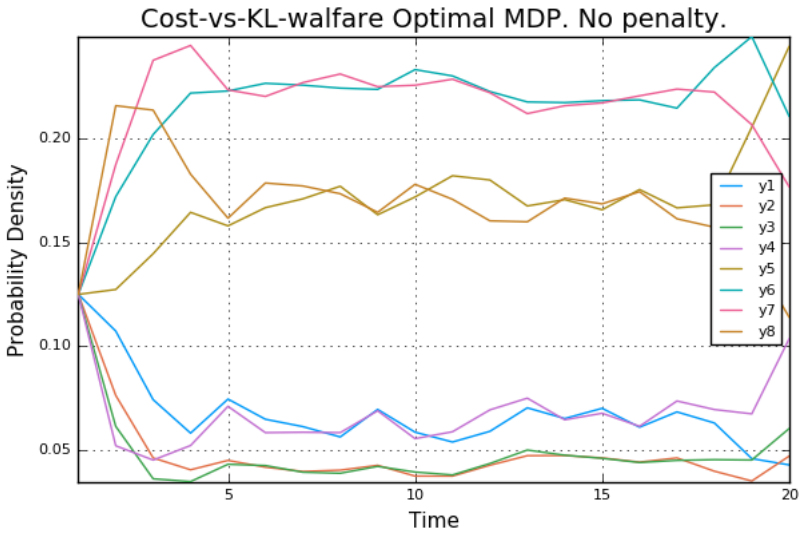

Fig. 1.2 Exemplary solution of the MDP problem for the case without penalty, i.e. with the objective represented by Eq. (1.4). The initial probability distribution, corresponds to the steady state of the "target" MC with the transition probabilities shown in Fig. (1.1). Eight curves show optimal dynamics of the respective, $\rho_\alpha(t), \quad \alpha = 1, \cdots, 8$.

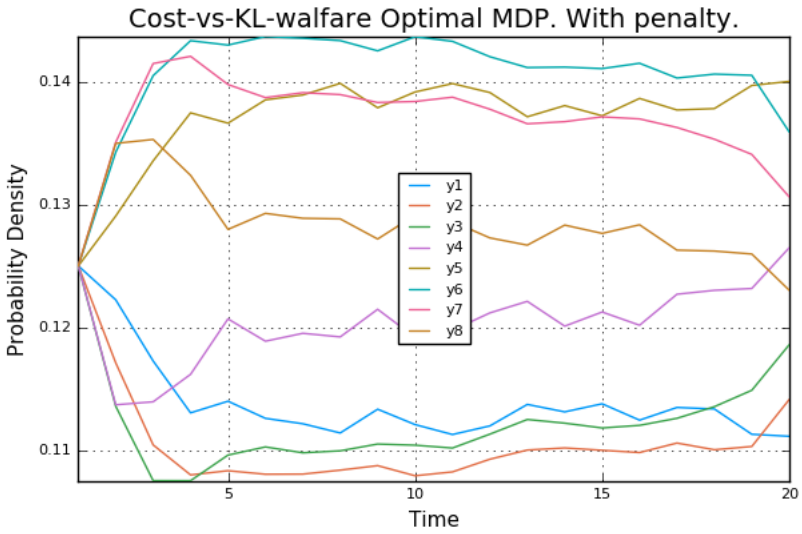

Fig. 1.3 Exemplary solution of the MDP problem for the case with penalty, i.e. with the objective represented by Eq. (1.6). The initial probability distribution, corresponds to the steady state of the "target" MC with the transition probabilities shown in Fig. (1.1). Eight curves show optimal dynamics of the respective, $\rho_\alpha(t), \quad \alpha = 1, \cdots, 8$ for the setting, equivalent to the one used in Fig. (1.2).

the experiments the exemplary case of 8 states and the target transition probability, $\bar{p}$, in the form shown in Fig. (1.1). For the same target, $\bar{p}$, we show the solution of the optimization problem that corresponds to the setting of Eq. (1.4) and Eq. (1.6), respectively, where in the latter case we choose a time-independent penalty factor $\gamma_{\alpha\beta} = 10$ for all transition except for those that correspond to advancing one (counter-clock-wise) step along the cycle. The special "along the cycle transitions" are not penalized, i.e. the respective transition factors are chosen equal to unity, $\gamma_{\mathrm{mod}[\alpha+1],\alpha} = 1, \quad \alpha = 1,\cdots,7$. We used the Algorithms described at the end of Section 1.2 and at the end of Section 1.3 respectively.

The results are shown in Figs. (1.2,1.3). Comparing the Figures, corresponding to the two regimes (with and without penalty), one observes that imposing penalty leads to a more homogeneous distribution ρ_α over the states, α.

1.5 Service (frequency control) Optimal

Another viable business model for an aggregator of an ensemble of cycling devices is to provide ancillary frequency control services to a regional SO. The auxiliary services consist in adjusting the energy consumption of the ensemble to an exogenous signal. Tracking the signal may require from some (or all) participants of the ensemble to sacrifice their natural cycling behavior. In this Section we aim to study, whether a perfect tracking of a pre-defined (exogenous) is feasible and then, in the case of feasibility, we would like to find the optimal solution causing least discomfort to the ensemble. Putting it formally, the aggregator solves the following optimization problem

$$\min_p \mathbb{E}_\rho \left[\underbrace{\sum_{t=0}^{T-1} \sum_{\alpha.\beta} \gamma_{\alpha\beta}(t) \log \frac{p_{\alpha\beta}(t)}{\bar{p}_{\alpha\beta}}}_{\text{welfare penalty}} \right] \tag{1.7}$$

$$\text{s.t.} \qquad \text{Eqs. (1.1)}$$

$$\underbrace{s(\tau) = \sum_\alpha \varepsilon_\alpha \rho_\alpha(\tau), \quad \forall\alpha, \quad \forall\tau = 1,\cdots,T}_{\text{energy tracking constraint}} \tag{1.8}$$

where $\bar{p}$ is the 'target' distribution represented by a stochastic matrix, e.g. leading to mixing into the standard steady state when $U = 0$; p is the stochastic, i.e., constrained by Eqs. (1.1), matrix; $s(t)$ is the amount of energy, requested by the system operator to balance the (transmission level) grid; ε_α is the amount of energy consumed by a device when it stays in the state α for a unit time slot. As stated above, we consider the most challenging setting which assumes that the tracking is per-

fect, that is the total consumption of the ensemble is exactly equal to the amount requested by the SO.

Introducing the Lagrangian multiplier for the energy tracking constraint, one restates Eq. (1.7) as the following min-max optimization

$$\min_{p,\rho}\max_{\xi}\left(\mathbb{E}_{\rho}\left[\sum_{t=0}^{T-1}\sum_{\alpha,\beta}\gamma_{\alpha\beta}(t)\log\frac{p_{\alpha\beta}(t)}{\bar{p}_{\alpha\beta}}+\sum_{t=1}^{T}\sum_{\alpha}\xi(t)\varepsilon_{\alpha}\right]-\sum_{t=1}^{T}\xi(t)s(t)\right)_{\text{Eqs. (1.1)}} \tag{1.9}$$

Reversing the order of summation in Eq. (1.9), we observe that optimization over p and ρ in the resulting expression becomes equivalent, under a substitution,

$$U_{\alpha}(\tau)=\xi(\tau)\varepsilon_{\alpha},\quad \forall\alpha,\quad \tau=1,\cdots,T. \tag{1.10}$$

to the penalized Kublack-Leibler welfare penalty optimization (1.6), discussed in the preceding Section 1.3.

In spite of the close relation between the energy tracking problem (1.8) and the profit optimality problem (1.6) the former is more difficult to implement. Indeed, while one can solve Eq. (1.6) in only one backward-forward run, we are not aware of existence of such an efficient algorithm for solving Eq. (1.8). The difficulty is related to the fact that $\xi(t)$, itself, should be derived in the result of a KKT condition, reenforcing the energy tracking constraint (1.8). A natural resolution of this problem is through an outer loop iteration including the following two sub-steps:

- Run the backward-forward penalty optimization algorithm described at the end of Section 1.3 using the current $\xi(t)$ (outer-step specific) profile.
- Update current $\xi(t)$ according to (1.8), utilizing the current $\rho(t)$ derived on the previous sub-step.

The outer-loop iterative scheme is initiated with $\xi(t)$ derived from (1.8) with $\rho(t)$ correspondent to the "normal" MP, $p\to\bar{p}$. One runs iterations till a pre-set tolerance is achieved. We plan to experiment with and analyze convergence of the iterative scheme in the future.

1.6 Hybrid Modeling: Towards Voltage Aware & Hierarchical Ensemble Control

In this Section we describe one possible scheme of an MDP approach integration into control of power distribution. The description here is sketchy and preliminary, meant to set up a formulation for further exploration.

The aggregation of many loads discussed in the manuscript so far has ignored the details of power flows as well as related voltage and Power Flow (PF) constrains. We have assumed that any of the discussed solutions is PF-feasible, in the sense that solution of PF equations are realizable for any of the consumption configuration and, moreover, the voltages and power flows along the power lines stay within the

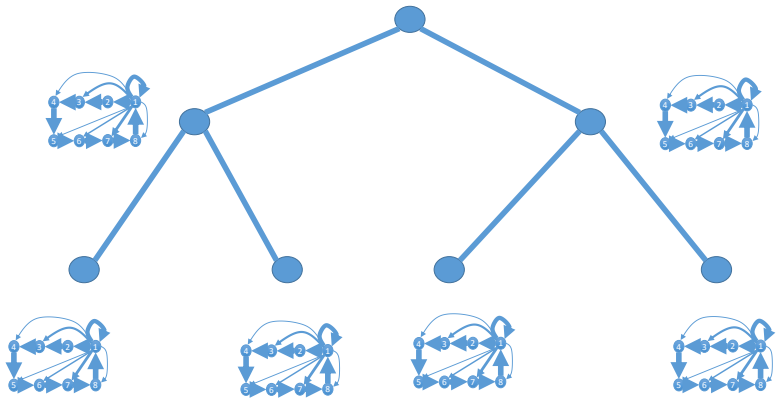

Fig. 1.4 Scheme illustrating hybrid/integrated and power-grid-aware MDP formulation. Nodes of the distribution level power system each is characterized by a sub-ensemble, modeled as a MP. Nodes are connected into a tree-like power system.

prescribed bounds. For these assumptions to hold, the aggregated devices should either be in the immediate geographical proximity, or otherwise the power-system should be operated with a significant safety margin. In general, to satisfy either of the two conditions globally and for a large ensemble is impractical. This motivates a discussion of the following hybrid model, stated in terms of many geographically-localized sub-ensembles connected into a power distribution network:

$$\min_{p,\rho,s,V} \sum_{t=0}^{T-1} \sum_i \sum_\alpha \left(U_\alpha^{(i)}(t+1) + \sum_\beta \gamma_{\alpha\beta}^{(i)}(t) \log \frac{p_{\alpha\beta}^{(i)}(t)}{\bar{p}_{\alpha\beta}^{(i)}} \right) \rho_\alpha^{(i)}(t) \tag{1.11}$$

$$s.t.$$

$$\rho_\alpha^{(i)}(t+1) = \sum_\beta p_{\alpha\beta}^{(i)}(t) \rho_\alpha^{(i)}(t), \quad \forall t, \quad \forall i, \quad \forall \alpha \tag{1.12}$$

$$\underbrace{s^{(i)}(t) \doteq \sum_\alpha \varepsilon_\alpha^{(i)} \rho_\alpha^{(i)} = V^{(i)}(t) \sum_{j\sim i} \left(\frac{V^{(i)}(t) - V^{(j)}(t)}{z_{ij}} \right)^*}_{\text{Nodal PF relations}}, \quad \forall i, \quad \forall t \tag{1.13}$$

$$\underbrace{\underline{v}^{(i)} \le |V^{(i)}|(t) \le \bar{v}^{(i)}}_{\text{voltage constraints}}, \quad \forall i, \quad \forall t. \tag{1.14}$$

where each sub-ensemble, representing, for example, stochastic/fluctuating consumption of a large apartment complex, is modeled as an aggregated MP. The objective (1.11) generalizes Eq. (1.6), thus accounting for the cost of services and

welfare quality (deviation from normal) for multiple ensembles. As earlier in the manuscript, we introduce the additional penalty factor, $\gamma_\alpha^{(i)}(t)$, in front of the KL (welfare penalty) term. This allows to reflect the significance of different ensembles, transitions, and times. Eqs.(1.12) generalize the ME (1.2) to sub-ensembles. Eqs. (1.13,1.14) introduce PF and voltage constraints into the MDP setting. One seeks for an optimum (1.11) over the vector $p = \left(p_{\alpha\beta}^{(i)}(t)\,|\forall i,\ \forall\alpha,\beta,\ \forall t\right)$ constructed from the stochastic transition-probability matrices, each representing a node in the network at all moments of time, over the vector $\rho = \left(\rho_\alpha^{(i)}(t)\,|\forall i,\ \forall\alpha,\ \forall t\right)$, reconstructed from p according to Eqs. (1.12), with the proper initial conditions provided, over the vector of voltages, $V = \left(V^{(i)}(t)\,|\forall i,\ \forall t\right)$.

Notice that other objectives, such as a contribution enforcing minimization of power losses in the distribution systems, other constraints, such as imposing bounds on line flows, as well as other controls, such as position of (voltage) tap-changers and/or reactive consumption at the nodes containing inverters, can be incorporated into the scheme in the spirit of [37, 34].

Since the MDP problem is stochastic by nature, it also allows incorporation of other stochastic sources, e.g. solar or wind renewables, which can be done either via modeling the stochastic sources as MP (with no control), or extending the modeling adding the so-called Chance-Constrained description in the spirit of [5].

An efficient solution of the hybrid problem (1.11) can be built via combining the techniques of temporal DP, developed in this manuscript for individual MDPs, with spatial (tree-graph) DP techniques, developed recently for OPF in power distributions [15] and taking advantage of the tree-graph operational layout of power distribution. We plan to work on practical implementation of these and other components of the hybrid model in the future.

1.7 Conclusions & Path Forward

In this manuscript we review and develop a computationally scalable approach for optimization of an ensemble of devices modeled via finite-space, finite-time MP. This approach builds upon earlier publications [20, 30, 31], addressing DR applications in power systems and beyond. A particularly practical and popular example of MP, relevant for power systems, is an ensemble of cycling loads, such as air-conditioners, water heaters, or residential water pumps [6, 33].

We have developed a number of useful MDP formulations aiming to achieve optimality for a diverse set of objectives of interest for an aggregator (of the ensemble) under different circumstances. We started the manuscript by describing MDP that balances overall expenses of the ensemble, acquired, when the cost of electricity varies in time, with welfare penalty that measures ensemble operational deviation from its normal behavior. Then we proceeded to discuss MDP, built to test feasibility of the ensemble to provide ancillary frequency control services to SO. Finally, we have addressed the future challenge of building a hybrid model that incorporates

MP and MDP modeling into deterministic and stochastic frameworks of Optimum Power Flows (operational dispatch) for power distribution.

The most important technical achievement of this paper is the development of an appropriate DP framework for stochastic optimization including DR. For the cost-vs-welfare optimization we have also tested the developed methodology and algorithms numerically (on a small scale example).

This manuscript did not offer a concluding fine-tuned summary of a completed project. Instead, we have focused here on presenting a new open-ended paradigm. In other words, we expect to see emergence of many more future extensions and generalizations of the approach which we have started to develop here. Some of these proposed future developments have been already discussed in the preceding Sections, especially in Section 1.6, others are briefly highlighted below:

- *Utilizing and Extending lin-solvability.* Intuitively, it is clear that lin-solvability is a rich property which should be advantageous both for analysis and algorithms. However, the use of this strong property in this manuscript was rather limited. We expect that the lin-solvability will bring an actual computational/algorithmic benefit not just for solving MDP per-se, but rather for solving more complex multi-level optimization/control problems, where an individual MDP represents an element of a richer model. (Some examples are mentioned below.) On the other hand, the lin-solvability described in Appendix 1.9.1.1 is a rather delicate property of the model, which is also easy to loose either partially or completely, as shown to be the case for the setting/formulation discussed in Appendix 1.9.1 and the main part of Appendix 1.9, respectively. This observation motivates further investigation of other settings/formulations, amenable for either full or partial lin-solvability. One promising direction for future analysis is to consider a generalization from the KL-divergence to Renyi-divergence, building upon and extending the approach of [16, 17].
- *Model Reduction.* Many models of power systems are too large and detailed for computations. This applies even to the routine Power Flow (PF) computation, which is a routing subtask for many power system problems of high level, involving optimization, control and generalizations accounting for stochasticity and robustness. Seeking to represent large-scale, long-time behavior, one is interested to build a reduced model, aggregating the small-scale and/or short-time details in a compact way. The reduction may be lossless or lossy, with or without ability to reconstruct the small-scale/short-time details. The developed here intrinsically stochastic MP framework is appropriate for the lossy case, where stochasticity represents the loss of insignificant details. We envision building reduced models, capable of more efficient, but still accurate computations, of say PF, in the format, similar to one, discussed in Section 1.6. The reduced model may consist of power lines with effective impedances and stochastic load/generation elements, represented by MPs. The level of coarse-graining may be pre-defined by a geography-preserving procedure, e.g., of the type discussed in [21], but it may also be left flexible/adaptive, where the number of states and allowable transitions for an individual MP is subject of optimization.

- *Hierarchical Control.* When the reduced model with MPs, representing aggregated loads/generations, is to be used in the context of optimization and control, changing from MP to MDP may be understood as allowing to split the optimization/control into two levels. Optimization with many MDPs modeling aggregated end-users, as discussed in Section 1.6, will produce optimal transition probability matrices for each individual MDP. Then, the task of implementing this policy is delegated to aggregators, responsible for individual MDPs. In this case, implementation means the global optimality/control is substituted by a two level hierarchical control. The scheme may be extended to represent more levels of control.
- *Derivation of MDP from a detailed Microscopic Model.* Consider an ensemble of continuous-time TCLs, each represented by a microscopic state in terms of temperature (continuous variable) and switch on/off status (discrete variables). In the case of an inhomogeneous ensemble, where each TCL or a group of TCLs may be parameterized differently, e.g., in terms of the allowed temperature rates and/or transition rates between on/off states, one is posing the question of representing this inhomogeneous ensemble as a single discrete time and discrete space (binned) MP or MDP. Developing methodology for constructing a representative MP/MDP for an aggregated ensemble is an important task for future research.
- *Accounting for Risk-Metric.* The MP/MDP methodology is sufficiently rich and flexible to account and mitigate for risks of different types, e.g., of overloading. Indeed, incorporating into MDP optimization additional constraints, limiting some elements of ρ or p, is a straightforward way of limiting the risk in probabilistic terms, natural for MDP.
- *Stochastic MDP.* MDP models already account for intrinsic stochasticity of an ensemble of devices (or coarse-grained areas), a model represents. However exogenous effects, such as those, representing the cost of electricity or the frequency signal, are modeled in this manuscript deterministically, even thought it is often more appropriate to represent generically uncertain and stochastic exogenous signal probabilistically. Formulating appropriate "second-order" statistical models represent an interesting challenge for future research.
- *From MDP to Reinforcement Learning (RL).* MDP is a standard tool, used in the field of RL. One of data-driven RL approaches [36], relevant for our discussion, suggests to consider "ideal" transition probabilities, $\bar{p}$, as unknown/uncertain, and then attempt to learn $\bar{p}$ from the data in parallel with solving the MDP. The approach was coined in [36] Z-learning and it is also closely related to the so-called Approximate DP (see [1] and references therein for many books and reviews). We anticipate that this general approach, applied to the models, introduced in the manuscript, will allow to build a data-driven and MDP-based framework for controlling an ensemble, whose normal behavior is known only through limited samples of representative behavior.
- *MDP for Supervised Learning (SL).* Recently a SL methodology was applied as a real-time proxy to solve difficult power system problems, such as to find an efficient description of the feasibility domain [18], to solve OPF [8] or power system reliability management problem [14]. The main idea in this line of research

rests on replacing an expensive power system computation with a ML black box, trained to evaluate a sufficient number of samples, labeled by relevant output characteristics. One interesting option consists in using properly designed MPs and/or MDPs for opening up the black box and turning it into (at least partially) physics-informed ML. Another possible direction would be to use MPs/MDPs as labels.

- *MDP as an element of a Graphical Model (GM) framework.* It was argued recently in [15, 10] that the approach of GM offers a flexible framework and efficient solutions/algorithms for a variety of optimization and control problems in energy systems (power systems and beyond). It is of interest to extend the GM framework, and build into it MDP methodology, formulations, and solutions.
- *Applications to Systems, e.g. Energy Systems.* The MDP models, discussed in this manuscript, are of interest well beyond representing ensembles (and coarse regions) in power systems. Similar methods and approaches are relevant for problems representing behavior and DR capabilities of consumers/producers' ensemble in other energy infrastructures, such as natural gas systems and district heating systems. In fact, the models discussed above fit practically "as is" to describe aggregation of many consumers of district heating systems, residing in a big apartment complex or densely populated residential area. The language of MDP is also universal enough to optimize joint energy consumption through multiple energy infrastructures of a "residential" ensemble (including electricity, heat and gas – possibly used as an alternative to central heating in local boilers).

1.8 Acknowledgments

The authors are grateful to S. Backhaus, I. Hiskens and D. Calloway for fruitful discussions, guidance, and valuable comments. The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy under Contract No. DE-AC52-06NA25396.

1.9 Dynamic/Bellman Programming

In this Appendix we discuss Dynamic Programming (DP) solution for the most general of the formulations, considered in this paper, for the profit-vs-welfare optimal formulations. Specifically we consider Eq. (1.6) introducing inhomogeneity in the inter-state transitions, expressed through the states and time dependent $\gamma_{\beta\alpha}(\tau)$ factors.

Let us restate Eq. (1.6) as

$$\min_{p} C\left(p|_0^{T-1}, \rho(0)\right)\Big|_{\sum_\beta p_{\alpha\beta}(\tau)=1,\ \forall\alpha}, \tag{1.15}$$

where $p|_0^{T-1}$ is the shortcut notation for the vector, constructed 0f the transition probability matrices evaluated at $t = 0, \cdots, T-1$, i.e., $(p(t)|\forall t = 0, \cdots, T-1$. The so-called value function in Eq. (1.15) can be decomposed according to

$$C\left(p|_0^{T-1}, \rho(0)\right) = C\left(p|_0^{\tau-1}, \rho(0)\right) + \sum_\alpha \varphi_\alpha\left(\tau, p|_\tau^{T-1}\right)\rho_\alpha\left(\tau, p|_0^{\tau-1}\right), \quad (1.16)$$

where the notation, $\rho_\alpha\left(\tau, p|_0^{\tau-1}\right)$ is introduced temporarily (just for the purpose of this derivation) to emphasize that ρ_α, computed at the moment of time τ, also depends, according to Eq. (1.2), on the transition probability matrices computed at all the preceding times. Here in Eq. (1.16) φ_α is defined at the final moment of time according to

$$\varphi_\alpha(T) = U_\alpha(T), \quad (1.17)$$

and then solved backwards in time by the following recursive equations

$$\forall\alpha, \quad \tau = T-1, \cdots 0: \quad \varphi_\alpha\left(\tau, p|_\tau^{T-1}\right) = \sum_\beta \varphi_\beta\left(\tau+1, p|_{\tau+1}^{T-1}\right) p_{\beta\alpha}(\tau)$$
$$+\sum_\beta \gamma_{\beta\alpha}(\tau) p_{\beta\alpha}(\tau)\log\left(\frac{p_{\beta\alpha}(\tau)}{\bar{p}_{\beta\alpha}}\right) + U_\alpha(\tau), \quad (1.18)$$

where, as before for p, the notation $\varphi_\alpha\left(\tau, p|_\tau^{T-1}\right)$ emphasizes that (by construction) $\varphi_\alpha(\tau)$ depends only on $p|_\tau^{T-1}$.

The dynamic programming-decomposed (recursive) structure of Eqs. (1.16,1.18) allows to evaluate optimization over p in Eq. (1.15) greedily as KarushKuhnTucker (KKT) first-order conditions - first over $p(T-1)$, then over $p(T-2)$, etc., all the way to $p(0)$, thus arriving at the following DP relations for the optimal p:

$$p_{\alpha\beta}(\tau) = \bar{p}_{\alpha\beta}\exp\left(-\frac{\varphi_\alpha(\tau+1) - \lambda_\beta(\tau)}{\gamma_{\alpha\beta}(\tau)}\right), \ \forall\alpha,\beta, \ \forall\tau = T-1, \cdots, 0, \quad (1.19)$$

where the Lagrange multipliers, $\lambda(\tau)$, are determined implicitly from the normalization conditions

$$\sum_\alpha \bar{p}_{\alpha\beta}\exp\left(-\frac{\varphi_\alpha(\tau+1) - \lambda_\beta(\tau)}{\gamma_{\alpha\beta}(\tau)}\right) = 1, \ \forall\beta, \ \forall\tau = T-1, \cdots, 0, \quad (1.20)$$

following from Eqs. (1.1). The Lagrange multipliers λ can also be derived via one-dimensional convex optimizations

$$\lambda_\beta(\tau) = \arg\min_\mu\left(\sum_\alpha \bar{p}_{\alpha\beta}\gamma_{\alpha\beta}(\tau)\exp\left(-\frac{\varphi_\alpha(\tau+1) - \mu}{\gamma_{\alpha\beta}(\tau)}\right) - \mu\right), \quad (1.21)$$
$$\forall\beta, \ \forall\tau = T-1, \cdots, 0.$$

Computationally, one may solve Eqs. (1.21) via gradient descent

$$\lambda_\beta(\tau) \leftarrow \lambda_\beta(\tau) - \delta\left(\sum_\alpha \bar{p}_{\alpha\beta} \exp\left(-\frac{\varphi_\alpha(\tau+1)-\lambda_\beta(\tau)}{\gamma_{\alpha\beta}(\tau)}\right)-1\right), \quad (1.22)$$
$$\forall\beta,\ \forall\tau = T-1,\cdots,0.$$

choosing an appropriate step, δ, and iterating Eq. (1.22), till the target tolerance of solution's accuracy is reached.

1.9.1 The case which allows explicit normalization

In the case of a general $\gamma(\tau)$ Eqs. (1.19) the Lagrange multipliers, $\lambda(t)$, cannot be expressed via p in a closed form. An exception is the case when $\gamma_{\alpha\beta}(\tau)$ is independent of α, i.e.

$$\gamma_{\alpha\beta}(\tau) \Rightarrow \gamma_\beta(\tau). \quad (1.23)$$

Then, Eq. (1.19) results in

$$p_{\beta\alpha}(\tau) = \frac{\exp\left(-\frac{\varphi_\beta(\tau+1)}{\gamma_\alpha(\tau)}\right)\bar{p}_{\beta\alpha}}{\sum_\nu \exp\left(-\frac{\varphi_\nu(\tau+1)}{\gamma_\alpha}\right)\bar{p}_{\nu\alpha}}, \quad \forall\alpha,\beta, \quad \forall t. \quad (1.24)$$

Substituting Eqs. (1.24) into Eqs. (1.18) one derives

$$\varphi_\alpha(\tau) = -\gamma_\alpha(\tau)\log\left(\sum_\beta \exp\left(-\frac{\varphi_\beta(\tau+1)}{\gamma_\alpha(\tau)}\right)\bar{p}_{\beta\alpha}\right) + U_\alpha(\tau), \quad \forall\alpha, \quad \forall\tau. \quad (1.25)$$

Notice that there are also other special cases which may allow for analytic expressions for the normalization. In particular, the case when $\gamma_{\alpha\beta}(\tau), \forall\tau, \quad \forall\alpha,\beta$ take values from a finite alphabet. We leave discussion of this and other interesting cases to future studies.

1.9.1.1 Linearly solvable case

Reduction (1.23) allows to map Eq. (1.15) to solution of DP equations (1.25), which are, however, nonlinear. When an additional reduction is made, limiting Eq. (1.23) even further to the state independent case,

$$\gamma_{\alpha\beta}(\tau) \Rightarrow \gamma(\tau), \quad (1.26)$$

Eqs. (1.25) reduce to what was coined in [36, 16, 17] the linearly-solvable MDP

$$\exp\left(-\frac{\varphi_\alpha(\tau)}{\gamma(t)}\right) \doteq u_\alpha(\tau) = \sum_\beta u_\beta(\tau+1)\bar{p}_{\beta\alpha}\exp\left(-\frac{U_\alpha(\tau)}{\gamma(t)}\right). \qquad (1.27)$$

Eqs. (1.27) are solved backwards in time with the final condition,

$$u_\alpha(T) = \exp\left(-\frac{U_\alpha(T)}{\gamma(T)}\right), \quad \forall\alpha. \qquad (1.28)$$

Two remarks are in order. First, we note that Eq. (1.26) is of a practical interest when one may be interested to change the relative importance of the welfare reenforcement vs price balance in the optimization. Second, other linearly solvable cases, in addition to those described by Eq. (1.26), may exist. We postpone a more general discussion of a broader class of the linearly solvable cases, as well as their practical utility, to future publications.

References

1. Reinforcement Learning, howpublished = https://en.wikipedia.org/wiki/reinforcement_learning, note = Accessed: 2017-01-07
2. Study on the demand response potential for seawater desalination projects (2016). URL http://www.ercot.com/content/wcm/lists/89476/Demand_Response_Potential_for_Seawater_Desalination_Projects_11_18_2016.pdf
3. Aunedi, M., Kountouriotis, P.A., Calderon, J.E.O., Angeli, D., Strbac, G.: Economic and environmental benefits of dynamic demand in providing frequency regulation. IEEE Transactions on Smart Grid **4**(4), 2036–2048 (2013). DOI 10.1109/TSG.2013.2258047
4. Bashash, S., Fathy, H.K.: Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In: Proceedings of the 2011 American Control Conference, pp. 4546–4553 (2011). DOI 10.1109/ACC.2011.5990939
5. Bienstock, D., Chertkov, M., Harnett, S.: Chance-constrained optimal power flow: Risk-aware network control under uncertainty. SIAM Review **56**(3), 461–495 (2014). DOI 10.1137/130910312. URL http://dx.doi.org/10.1137/130910312
6. Callaway, D., Hiskens, I.: Achieving controllability of electric loads. Proceedings of the IEEE **99**(1), 184–199 (2011). DOI 10.1109/JPROC.2010.2081652
7. Callaway, D.S.: Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. Energy Conversion and Management **50**(5), 1389 – 1400 (2009). DOI http://dx.doi.org/10.1016/j.enconman.2008.12.012. URL http://www.sciencedirect.com/science/article/pii/S0196890408004780
8. Canyasse, R., Dalal, G., Mannor, S.: Supervised Learning for Optimal Power Flow as a Real-Time Proxy. ArXiv e-prints (2016)
9. Chertkov, M., Chernyak, V.: Ensemble of thermostatically controlled loads: Statistical physics approach (2017)
10. Chertkov, M., Krishnamurthy, D., Misra, S., Hentenryck, P.V., Vuffray, M.: Graphical models and belief propagation-hierarchy for optimal physics-constrained network flows (2017)
11. Chong, C.Y., Debs, A.S.: Statistical synthesis of power system functional load models. In: Decision and Control including the Symposium on Adaptive Processes, 1979 18th IEEE Conference on, vol. 2, pp. 264–269 (1979). DOI 10.1109/CDC.1979.270177
12. Chong, C.Y., Malhami, R.P.: Statistical synthesis of physically based load models with applications to cold load pickup. Power Apparatus and Systems, IEEE Transactions on **PAS-103**(7), 1621–1628 (1984). DOI 10.1109/TPAS.1984.318643

13. Donnelly, M., Harvey, D., Munson, R., Trudnowski, D.: Frequency and stability control using decentralized intelligent loads: Benefits and pitfalls. In: IEEE PES General Meeting, pp. 1–6 (2010). DOI 10.1109/PES.2010.5589835
14. Duchesne, L.: Machine learning of proxies for power systems reliability management. Master's thesis, Faculty of Applied Sciences, Department of Electrical Engineering and Computer Science, University of Liege, Belgium (2016). URL https://matheo.ulg.ac.be/bitstream/2268.2/1374/4/master_thesis_laurine_duchesne.pdf
15. Dvijotham, K., Chertkov, M., Van Hentenryck, P., Vuffray, M., Misra, S.: Graphical models for optimal power flow. Constraints pp. 1–26 (2016). DOI 10.1007/s10601-016-9253-y. URL http://dx.doi.org/10.1007/s10601-016-9253-y
16. Dvijotham, K., Todorov, E.: A Unifying Framework for Linearly Solvable Control. ArXiv e-prints (2012)
17. Dvijotham, K., Todorov, E.: Linearly Solvable Optimal Control, pp. 119–141. John Wiley & Sons, Inc. (2013). DOI 10.1002/9781118453988.ch6. URL http://dx.doi.org/10.1002/9781118453988.ch6
18. Dymarsky, A., Horodnytskyi, O., Maximov, Y., Turitsyn, K.: Statistical learning of power flow feasibility domain (2017)
19. Fleming, W.H., Mitter, S.K.: Optimal control and nonlinear filtering for nondegenerate diffusion processes. Stochastics **8**(1), 63–77 (1982). DOI 10.1080/17442508208833228. URL http://dx.doi.org/10.1080/17442508208833228
20. Ghaffari, A., Moura, S., Krstic, M.: Pde-based modeling, control, and stability analysis of heterogeneous thermostatically controlled load populations. Journal of Dynamic Systems Measurement and Control **137**(10) (2015)
21. Grudzien, C., Deka, D., Chertkov, M., Backhaus, S.: Conceptual structure preserving reduction of large-scale power system models (2017)
22. Ihara, S., Schweppe, F.: Physically based modeling of cold load pickup. Power Apparatus and Systems, IEEE Transactions on **PAS-100**(9), 4142–4150 (1981). DOI 10.1109/TPAS.1981.316965
23. Kappen, H.J.: Linear theory for control of nonlinear stochastic systems. Phys. Rev. Lett. **95**, 200,201 (2005). DOI 10.1103/PhysRevLett.95.200201. URL http://link.aps.org/doi/10.1103/PhysRevLett.95.200201
24. Kontoyiannis, I., Meyn, S.: Large deviations asymptotics and the spectral theory of multiplicatively regular markov processes. Electron. J. Probab. **10**, no. 3, 61–123 (2005). DOI 10.1214/EJP.v10-231. URL http://ejp.ejpecp.org/article/view/231
25. Lu, N., Chassin, D.: A state-queueing model of thermostatically controlled appliances. Power Systems, IEEE Transactions on **19**(3), 1666–1673 (2004). DOI 10.1109/TPWRS.2004.831700
26. Lu, N., Chassin, D., Widergren, S.: Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model. Power Systems, IEEE Transactions on **20**(2), 725–733 (2005). DOI 10.1109/TPWRS.2005.846072
27. Malhame, R., Chong, C.Y.: Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. IEEE Transactions on Automatic Control **30**(9), 854–860 (1985). DOI 10.1109/TAC.1985.1104071
28. Malhame, R., Chong, C.Y.: On the statistical properties of a cyclic diffusion process arising in the modeling of thermostat-controlled electric power system loads. SIAM Journal on Applied Mathematics **48**(2), 465–480 (1988). DOI 10.1137/0148026. URL http://dx.doi.org/10.1137/0148026
29. Meyn, S., Barooah, P., Busic, A., Chen, Y., Ehren, J.: Ancillary service to the grid using intelligent deferrable loads. Automatic Control, IEEE Transactions on **60**(11), 2847–2862 (2015). DOI 10.1109/TAC.2015.2414772
30. Meyn, S.P., Barooah, P., Busic, A., Chen, Y., Ehren, J.: Ancillary service to the grid using intelligent deferrable loads. IEEE Transactions on Automatic Control **60**(11), 2847–2862 (2015). DOI 10.1109/TAC.2015.2414772

31. Paccagnan, D., Kamgarpour, M., Lygeros, J.: On the range of feasible power trajectories for a population of thermostatically controlled loads. In: 2015 54th IEEE Conference on Decision and Control (CDC), pp. 5883–5888 (2015). DOI 10.1109/CDC.2015.7403144
32. Schweppe, F.C., Tabors, R.D., Kirtley, J.L., Outhred, H.R., Pickel, F.H., Cox, A.J.: Homeostatic utility control. IEEE Transactions on Power Apparatus and Systems **PAS-99**(3), 1151–1163 (1980). DOI 10.1109/TPAS.1980.319745
33. Siano, P.: Demand response and smart gridsa survey. Renewable and Sustainable Energy Reviews **30**, 461 – 478 (2014). DOI http://dx.doi.org/10.1016/j.rser.2013.10.022. URL http://www.sciencedirect.com/science/article/pii/S1364032113007211
34. Sulc, P., Backhaus, S., Chertkov, M.: Optimal distributed control of reactive power via the alternating direction method of multipliers. IEEE Transactions on Energy Conversion **29** (2014). URL https://arxiv.org/abs/1310.5748
35. Todd, D., Caufield, M., Helms, B., Starke, M., Kirby, B., Kueck, J.: Providing reliability services through demand response: A preliminary evaluation of the demand response capabilities of alcoa inc. (2009). URL https://www.ferc.gov/eventcalendar/Files/20100526085850-ALCOA%20Study.pdf
36. Todorov, E.: Linearly-solvable markov decision problems. In: B. Schölkopf, J. Platt, T. Hoffman (eds.) Advances in Neural Information Processing Systems 19, pp. 1369–1376. MIT Press (2007). URL http://papers.nips.cc/paper/3002-linearly-solvable-markov-decision-problems.pdf
37. Turitsyn, K., Sulc, P., Backhaus, S., Chertkov, M.: Options for control of reactive power by distributed photovoltaic generators. Proceedings of the IEEE **99**(6), 1063–1073 (2011)
38. Zheng, Y., Hill, D.J., Zhang, C., Meng, K.: Non-interruptive thermostatically controlled load for primary frequency support. In: 2016 IEEE Power and Energy Society General Meeting (PESGM), pp. 1–5 (2016). DOI 10.1109/PESGM.2016.7741469